

TCP/IP NETWORK, IP CONFIGURATION, AND NAME RESOLUTION STRATEGIES

**After reading this chapter and completing the exercises,
you will be able to:**

- ◆ Design a TCP/IP networking strategy
- ◆ Design IP configuration strategies using static addressing and DHCP
- ◆ Design a name resolution strategy using DNS
- ◆ Design a name resolution strategy using WINS

You say you already know all you need to know about IP addressing and name resolution? That's great! Then you'll be all the more prepared to take advantage of the information in this chapter as we show you how Microsoft has made your job more interesting by adding new features and capabilities in these areas. In this chapter, you will be shown strategies for using public and private addresses and for taking a network address and subnetting it to fit your network. Then we'll move onto strategies for configuring TCP/IP computers through DHCP. Finally, you will work with naming strategies—using both DNS and WINS, where appropriate.

DESIGNING A TCP/IP NETWORK

When designing a TCP/IP network, smart administrators implement the latest in security and performance enhancements. Keeping up with the newest tips and techniques allows them to design the best networks that technology can “buy.” We discuss these in the following sections.

TCP/IP Security Features in Windows 2000

Internal security, which is the act of protecting your network against internal threats, is a current hot topic in network design. Several enhancements to TCP/IP in Windows 2000 target internal security.

One enhancement is **Application-layer packet filtering**, which allows filtering of packets on a host-by-host basis. When you turn on packet filtering, you block all packet types except those listed. This allows you to control what travels on your network. Further, you may block all packets to a certain **port** except those sent from a specified address. If you choose to include packet filtering in your design, be sure to test it with all the applications you expect to run so that you don’t disable a port that an application or service depends on for functionality. Figure 4-1 shows the enabling of filtering.

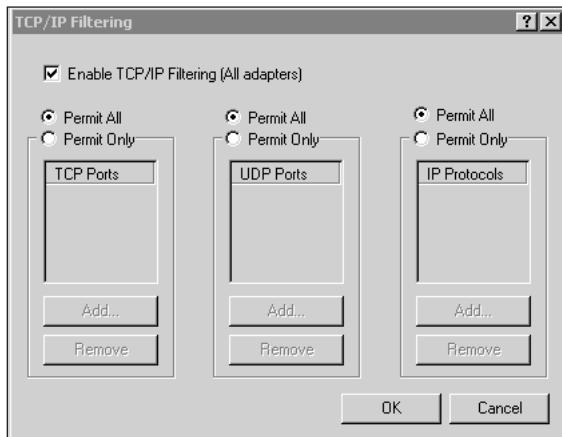


Figure 4-1 TCP/IP filtering

Protocols play a role in security as well. For instance, **IP Security (IPSec)** was developed by the IETF for the next version of IP, IPv6, and as an optional extension to IPv4. As such, Microsoft included it in Windows 2000. IPSec allows authentication of source and destination **hosts** before encryption of the data packets and before data is sent. When two machines that support IPSec establish a connection, they negotiate security settings; then all subsequent traffic over that connection is subject to the security settings. This is all transparent to the users and applications generating the traffic over the connection.

IPSec has two modes: tunnel mode and transport mode. In **tunnel mode**, IPSec will encapsulate IP packets and optionally encrypt them. A designer might use IPSec tunnel mode in a network design when traffic must use an untrusted network using routers or gateways that do not support virtual private networks (VPNs) using Layer 2 Tunneling Protocol (L2TP) or Point-to-Point Tunneling Protocol (PPTP). In **transport mode**, IPSec can be used to authenticate and/or encrypt communications between computers without using a tunnel. This provides security for your intranet traffic without the overhead of VPNs.

There are several components used by IPSec.

- **IPSec Driver:** Handles the actual protocol-level tasks of encrypting, authenticating, verifying, and decrypting packets. This driver can offload some of the cryptographic processing from the host CPU to that of the NIC when the NIC supports on-board processing.
- **Authentication Header (AH):** A protocol of IPSec that digitally signs the contents of packets to protect against replay attacks, tampering, and spoofing. It does not ensure confidentiality.
- **Encapsulating Security Payload (ESP):** A protocol of IPSec that ensures data confidentiality by encrypting the entire contents of each packet. It can be used alone or in combination with AH for greater security.
- **Internet Key Exchange (IKE):** A protocol that generates the keys for the IPSec protocols and negotiates keys for other protocols that require keys.
- **Internet Security Association Key Management Protocol (ISAKMP):** An IPSec protocol that provides the method by which two computers can agree on a common set of security settings and a secure way for them to exchange a set of encryption keys to use for their communication.
- **Oakley:** A key determination protocol that uses the Diffie-Hellman key exchange algorithm.
- **Transforms:** Defines a set of actions (or transformations) that is applied to data for security purposes. This can include the algorithm for encryption, the key sizes and method for deriving them, and the process used to encrypt the data. An example of a transform is the DES-DBC transform used by ESP.
- **IPSec Policy Agent:** Runs as a service on Windows 2000 machines. When the operating system starts, this component retrieves the IPSec policy settings from Active Directory and applies them.

Administrators can configure the IPSec policy through local or group policy security settings for all computers that need elevated security while communicating. Then, when one of these computers initiates communications with another computer, the following occurs:

1. The IPSec driver and the ISAKMP receive the IPSec policy settings.
2. ISAKMP negotiates between hosts, based on their policy settings, and builds a **security association (SA)**.

3. The Oakley protocol is used to negotiate a master key that can be used to secure further IPSec negotiations. After this, the two machines can negotiate the actual IPSec settings to be used for the connection. A second SA comprises this set of security methods and the keys that are used.
4. Then, based on the security policy agreed on for the session, the IPSec driver monitors, filters, and secures the Transport layer against network traffic.

When you use a VPN, all your traffic has to go across it and you can only talk to computers that are using the VPN. When using IPSec transport mode, you can communicate with any computers that comply with your configured IPSec policy without requiring a tunnel. Further benefits include the fact that administrators have flexibility in configuring the method of authentication, data integrity, and data encryption based on a predefined security policy. In addition, IPSec authentication ensures data integrity using an AH.



Only use AH when you need data integrity without encryption or if encryption is provided by another component.

ESP is a component you can use in combination with AH to provide encryption, but nothing comes without a price. This combination of configuring IPSec to use ESP to provide both authentication and encryption can be very processor-intensive.



Microsoft's pre-Windows 2000 clients do not support IPSec, although third-party products will allow Windows 9x, Windows NT, and several flavors of UNIX to use IPSec. For a list of such products, go to the Web site of ICSA Labs at www.icsalabs.com/html/communities/ipssec/certification/certified_products/.

For computers to communicate using IPSec, the following must be true:

- An IPSec policy must exist, but only one can be active on a computer at any time.
- The IKE process must occur in order to arrive at the DA, the agreed-upon security level, and the keys to be used.
- During the session, all data exchange must occur per the SA.

In a Windows 2000 domain, you will want to define IPSec policy at the group policy level. Figure 4-2 shows the IPSec policy in group policy. If your Windows 2000 computers are not in a Windows 2000 domain, define IPSec policy at the local policy level.

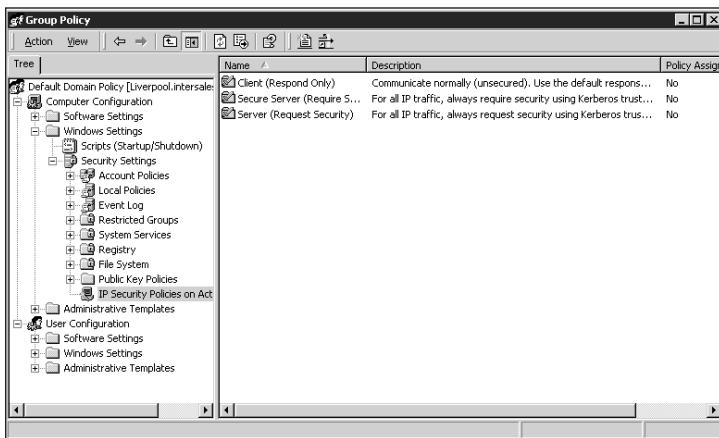


Figure 4-2 IPSec policy in Group Policy

Windows 2000 has three predefined IPSec policies that can be used “as is” or modified:

- **Client (Respond Only):** This policy should be used for computers that only use IPSec when establishing a session with another computer that requires IPSec. With this policy, the computer itself does not initiate IPSec usage, but is able to participate when another computer requires it.
- **Server (Request Security):** This policy provides a middle ground. Use it when the network includes hosts that are IPSec-enabled as well as computers that are not IPSec-enabled, and the host you are configuring must communicate with both. When enabled, the host will request the use of IPSec, but if the other host is not configured for IPSec, it will still allow communications.
- **Secure Server (Require Security):** This policy requires IPSec security and will allow no communication with hosts that cannot participate in IPSec communication. Use this when you require a high level of security and are sure that the server only needs to communicate with other IPSec-enabled hosts.

IPSec cannot be used through any process that translates addresses. If you need to use IPSec between a private network and a public network through a security gateway, firewall, proxy server, or routers that are performing traffic filtering, you need to open certain ports:

- For IPSec Authentication Header traffic, open IP Protocol ID 51.
- For IPSec Encapsulating Security Protocol traffic, open IP Protocol ID 50 and UDP Port 500.

Anytime you include security measures in your network design, you will have to balance security against performance. Remember that the higher the security, the greater the negative effect on performance.

TCP/IP Performance Enhancements in Windows 2000

Windows 2000 has new built-in performance enhancement features for TCP/IP. We discuss them in the following sections. You need to know about these enhancements because they will help you get the most out of your network design, allowing you to provide the performance required without additional costs.

Large TCP Windows and Selective Acknowledgment

With the flow-control method known as “sliding windows,” each host has two windows—a send window and a receive window. The size of the TCP receive window on each host is a function of the size of its receive buffer, which is the maximum amount of data the sending host can send before it must wait for an acknowledgment from the receiving host.

During initiation of the session, TCP on each host sets its Send window to the size of the receive window on the other host. The previous maximum to this window size was 64 KB, which has proven inadequate for the newer and faster WAN technologies, such as fiber optics. The problem here is that, although there is quite a bit of bandwidth, the return acknowledgment takes a long time. Therefore, the sending host will resend because it has not received an acknowledgment.

The solution to the problem, as implemented in Windows 2000, is twofold. First, it now supports a new TCP option called “window scale.” With it, the Windows 2000 TCP/IP stack uses larger default window sizes than previous versions. An administrator can also manually control this size by editing the registry on the Windows 2000 computers.



Microsoft Knowledge Base article Q224829 describes how to modify this setting, and article Q263088 provides further information.

The second part of the solution is based on the fact that recovery from segment loss is enhanced by the use of selective acknowledgment. The traditional TCP acknowledgment scheme has a problem: If a segment at the beginning of the send window is not received, but all other segments are, then none will be acknowledged until the missing segment is received. This will cause the sender to repeatedly retransmit segments that have already been received until it receives an acknowledgment of receipt of the missing segment. Selective acknowledgment (SACK) allows for the acknowledgment of receipt of non-contiguous segments. Thus the sender will resend only the missing segments.

ICMP Router Discovery

Another new feature is the fact that a computer running routing and remote access can be configured to perform **Internet Control Message Protocol (ICMP) router discovery**, by which a host can discover a router automatically, in spite of not having a default gateway configured in its TCP/IP properties. This is disabled by default on

Windows 2000 hosts, but can be configured by DHCP. In addition, a Windows 2000 server with routing and remote access enabled may also have this turned on.

Disabling NetBIOS over TCP/IP

Prior to Windows 2000, Microsoft network clients depended on NetBIOS to reference Microsoft network resources and use the NetBIOS namespace, so the NetBIOS Session-layer protocol could not be removed or disabled from a Windows computer with a Microsoft client.

Windows 2000 has removed this dependency for Active Directory domains and Windows 2000 clients. However, non-Windows 2000 Microsoft network clients still depend on NetBIOS, as do old applications. If you are absolutely sure you do not need it, you may disable NetBIOS through the Advanced button on the TCP/IP Properties dialog box. Clicking the bottom produces the dialog box in Figure 4-3.

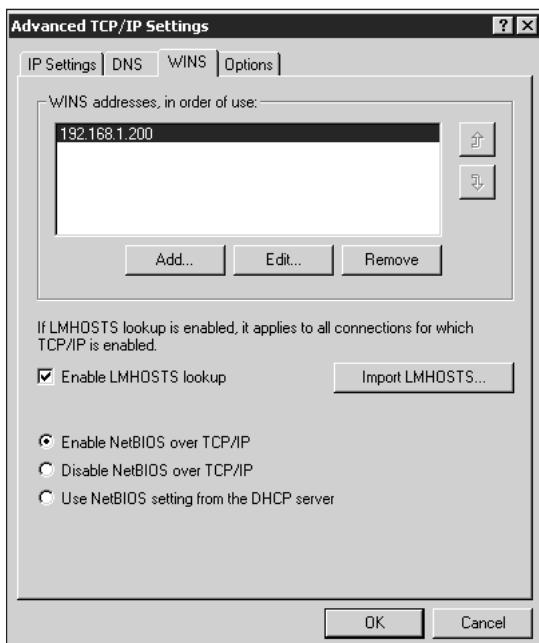


Figure 4-3 Advanced dialog box in TCP/IP Properties

If NetBIOS is disabled, the computer will no longer listen to NetBIOS traffic on that interface, and if it is a WINS server, it cannot respond to WINS client registration and query requests on that interface. Therefore, do not disable NetBIOS unless you are certain that there is no dependency on it.

As a rule, you will only disable NetBIOS under special circumstances and on certain hosts. One instance is when a computer is multi-homed, having one network interface

on a private network and another on a public network. You could disable NetBIOS on the external interface. Another circumstance is on edge proxy servers or **bastion hosts** (a gateway between an inside network and an outside network designed to defend against attacks aimed at the inside network) used behind a firewall and/or within a screened subnet. A screened subnet is a subnet between firewalls; it is also known, somewhat facetiously, as a **demilitarized zone (DMZ)**.

Quality of Service (QoS)

Jitter is the period frequency displacement of the signal from its ideal location. In other words, it is the variation in the delay in your transmission. This poses unique issues for real-time traffic, particularly voice- or mission-critical data transfers.

Delays—or jitter—are caused by different situations and entities on the network. For instance, consider WAN links. Further consider a router with two Fast Ethernet interfaces, which, as you know, are theoretically 100 Mbps, and a serial interface connected to a T-1 that is operating at 1.544 Mbps. In this case, it would not be unusual for that router to receive much more traffic on its LAN interfaces than it could transmit across the WAN. And when the router’s buffers fill up, it is forced to start dropping packets. When it drops a packet, it can send an Internet Control Message Protocol (ICMP) Source Quench message to the sender, which instructs it to reduce its transmit rate. Although this works great for normal data traffic, it is terrible for real-time traffic. First of all, real-time traffic usually is connectionless, which means, when packets are dropped, they’re gone forever and never re-sent. Second, the packet has to hang out in the buffer until all the other packets in front of it have been transmitted.

The solution to this problem is **Quality of Service (QoS)** and its two main functions: prioritization and resource reservation. Prioritization works like this: The network hardware manufacturers have replaced their single buffers with multiple buffers or “queues.” These queues are assigned percentages of bandwidth. So a router, for instance, may take a couple of packets out of the “high priority” queue, and then one packet out of the “medium priority” queue, and repeat this process as often as needed. This process lets high-priority traffic get preference over low-priority traffic, but how does the router know which buffer to place a packet into? Well, that’s where the IP precedence field comes in. In the header of every IP packet is a three-bit field that contains a number from zero to seven. The source of the packet simply assigns a higher number to important packets. The router administrator configures the routers to know which number goes into which queue.

Resource reservation is generally implemented with the Resource Reservations Protocol (RSVP), which allows an RSVP-aware application to request dedicated bandwidth from the network devices. The routers receive the request and check to see if they have the bandwidth available and then accept or deny the request. Once accepted, the host is almost guaranteed sufficient bandwidth to transmit its information.

Both prioritization and RSVP are ideally implemented in conjunction with a policy management system. This system includes Policy Enforcement Points (like routers and firewalls) and Policy Decision Points (like policy servers, which contain rules or policies). All of this is tied together with a directory, like Active Directory, which allows all the devices in the network to work together to get the important traffic where it needs to go. Microsoft has begun to implement all of this in a set of components collectively called QoS. The following are components involved in the Windows 2000 QoS story:

- **Generic QoS Application Programming Interface (GQoS API):** This is an Application Programming Interface through which programmers can add the ability to specify or request bandwidth to their applications.
- **Resource Reservations Protocol (RSVP):** The media-independent signaling component of QoS which establishes end-to-end communications. RSVP is implemented as RSVP.EXE.
- **Resource Reservations Protocol Service Provider (RSVP SP):** The service provider that accesses RSVP.EXE to initiate the establishment of a reservation.
- **Traffic Control (TRAFFIC.DLL):** It uses the parameters defined for the QoS communication to regulate traffic. The GQoS API calls up TRAFFIC.DLL.
- **Generic Packet Classifier (MSGPC.SYS):** This component determines the class of the packet.
- **QoS Packet Scheduler (PSCHED.SYS):** This component enforces the traffic flow parameters.
- **QoS Admission Control Service (QoS ACS):** This component serves as a central clearinghouse, allowing or denying bandwidth requests. In a high-traffic situation, you may want to plan to have this service available on each segment of the network.
- **Local Policy Module (MSIDLPM.DLL):** This is the policy-enforcement and policy-decision point of QoS. It is used by the QoS ACS to view the user name in the RSVP message and to determine if this user is permitted to use QoS per the admission control policy in Active Directory.

The QoS Admission Control host can be installed on a Windows 2000 server through Add/Remove Programs, Networking Services and is limited to controlling the subnet(s) to which it is directly connected. Once this is installed, the QoS Admission Control console can be accessed from the Administrative Tools menu. Figure 4-4 shows the console.

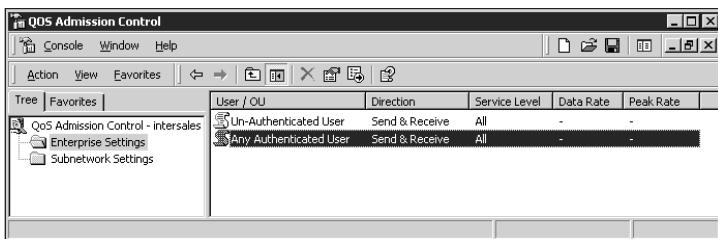


Figure 4-4 QoS Admission Control console

The components of QoS work in the following manner:

1. QoS-aware software requests QoS service, calling the RSVP SP.
2. RSVP SP, in turn, calls the RSVP service requesting the necessary bandwidth.
3. An RSVP message is sent to the QoS ACS server with a bandwidth request.
4. The QoS ACS checks that the available bandwidth can accommodate the request and calls the local policy module, through which it examines the user policy and compares it to the admission control policy in Active Directory.
5. If permissions were granted, QoS ACS allocates the bandwidth and passes the message to the receiving host.
6. QoS-enabled routers are configured with the bandwidth reservation, and wait for the receiving host to return the RSVP message.
7. As the receiver's message arrives at each router, the router decides whether to accept the reservation and commit bandwidth. If all intervening routers commit the bandwidth, the sender can begin the transmission.

Windows 2000 TCP/IP Redundancy Features

Windows 2000 shipped with two components that provided redundancy in the form of clustering. They are Network Load Balancing Service (NLB) and Microsoft Cluster Service (MSCS). The third component, which came out of Microsoft Application Server, is Component Services Load Balancing. We discuss the first two in the following sections; the third we will save for Chapter 9.

NLB is not completely new to Microsoft networking, since it was available for Windows NT 4.0 Server, Enterprise Edition for a few years before Windows 2000. However, it is now an integral part of both Windows 2000 Advanced Server and Datacenter Server, and will not require special service packs separate from the main product. This service allows up to 32 servers to appear to be a single server, with NLB balancing the load of incoming TCP/IP traffic among the servers. This is ideal for multiple servers, such as Web servers, hosting static data. Figure 4-5 is a simple example of NLB.

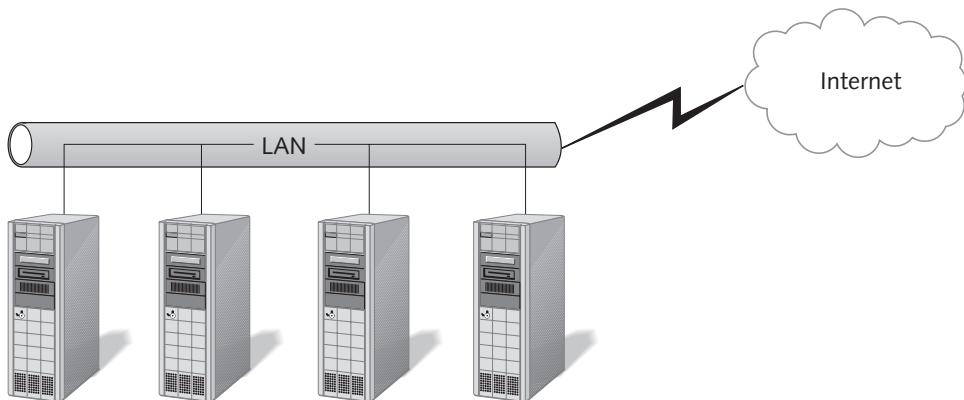


Figure 4-5 Network Load Balancing

MSCS, like NLB, is now integrated as a service into Windows 2000 Advanced Server and Windows 2000 Datacenter Server. Windows NT 4.0 Server, Enterprise Edition also offered clustering in a version called Windows Load Balancing Service (WLBS), so the service is not completely new to Windows 2000.

Windows 2000 Advanced Server supports a two-server failover, and Windows 2000 Datacenter Server supports a four-server, cascading failover. A **failover** occurs in a server cluster when one server automatically takes over for a failed server. An Advanced Server MSCS cluster would typically have two servers—one primary and the other secondary (which is constantly checking the status of the primary and which takes over if the primary fails). Each server has an internal drive on which the operating system and MSCS run. They would share a large SCSI disk or (more likely) disk array for the use of the clustered application or service. See Figure 4-6 for an example. Obviously, Datacenter Server, with a four-server, cascading failover, would be most appropriate for a mission-critical application. Several applications, such as Exchange 5.5, Enterprise Edition, and SQL Server 7.0 and the “2000” versions of these applications, are cluster-aware. There are also several network services, such as DHCP, WINS, and DNS, that are cluster aware.

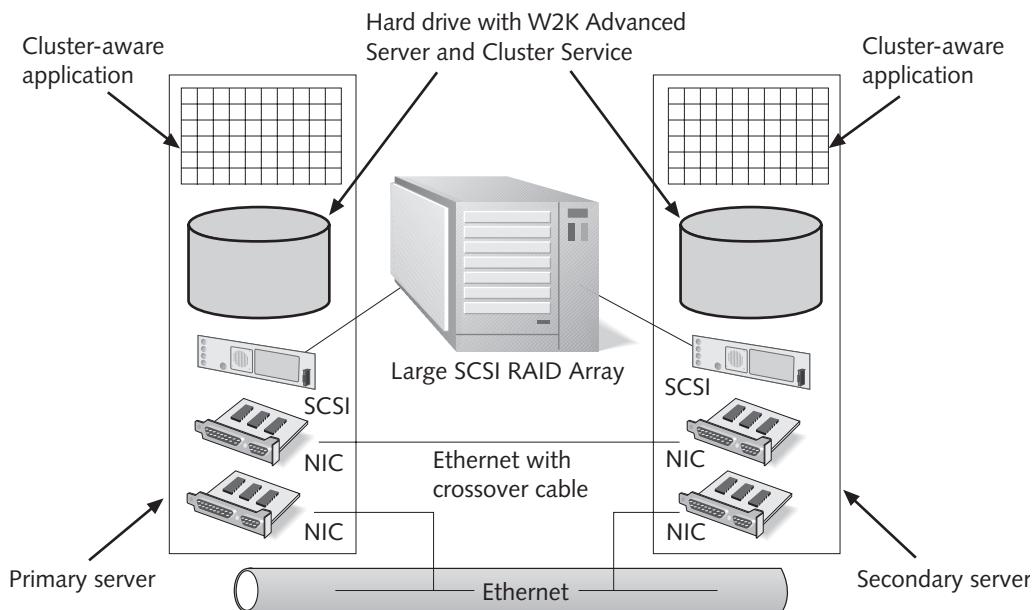


Figure 4-6 Microsoft Cluster Service

Designing TCP/IP Addressing and the Implementation Plan

As a network designer, you must devise a plan for using IP addresses effectively and efficiently. In this section, we explore the use of public addresses, private addresses, subnetting, supernetting, and variable-length subnet masks (VLSMs) in your design.



The following discussion assumes that you have a thorough understanding of TCP/IP. If you feel you need a review before proceeding, we suggest two Web sites for excellent tutorials on IP addressing and other TCP/IP topics: www.learntosubnet.com and www.ralphb.net/IPSubnet. We also recommend the following Course Technology books available at www.coursesdirect.com: *CCNA Guide to Cisco Networking Fundamentals* and *Microsoft Introduction to TCP/IP Internetworking*.

Obtaining Public Addresses

Before Internet addresses became so scarce, an organization would obtain its public address from a regional registrar. Now, however, with ISPs providing clients with addresses from portions of their address block, if you are a client, you are only “leasing” the address, not “owning” it. Fortunately, if you change ISPs, the domain name registration associated with the address is usually changed to show the new ISP’s name servers as being authoritative.



Most of us “normal mortals” have only heard of the Internet Assigned Numbers Authority (IANA), and believe that is where organizations and ISPs must go for addresses, but IANA is the overall number authority and the American Registry for Internet Numbers (ARIN) is one of several address registrars. Learn more about ARIN and the other regional registrars at www.arin.net.

Private Addresses

Private addresses, which are reserved for private internetworks, are never found in Internet routing tables. The blocks of private addresses specified through Internet standards are as follows:

- **10.0.0.0/8:** A class A private network address includes a range of addresses from 10.0.0.1 to 10.255.255.254. It has 24 host bits and provides the greatest number of subnet and host configurations.
- **172.16.0.0/12:** A range of 16 class B networks. It consists of all the addresses in the range of 172.16.0.1 to 172.31.255.254.
- **192.168.0.0/16:** A range of 256 class C network IDs, consisting of all the addresses in the range of 192.168.0.1 to 192.168.255.254.

These addresses will never be assigned as public addresses. Therefore, Internet routers are configured to never route these addresses. Hosts using these addresses on a private network that want to access resources on the Internet must go through a device that provides Network Address Translation (NAT) or uses an Application-layer gateway, such as a proxy server. In both cases, the host providing the address translation or gateway must have an interface with a valid Internet address.

Addressing Strategies

IP addressing strategies include subnetting, supernetting, and variable-length subnetting. In **subnetting**, you take a network address, such as 192.168.0.0/16, and “borrow” bits from the host portion to subdivide this single network address into multiple network addresses. This is at the core of your addressing strategies.

Supernetting is a strategy in which you borrow bits from the network portion to combine several network addresses into one. For example, if a company had 800 hosts, at one time they might have applied for a class B network ID. It is doubtful if they would have gotten it, but if they could, it would be a waste of thousands of host addresses. They could also have applied for four class C addresses, but this requires four entries in the Internet backbone router tables.

Rather than allowing the class B or class C solution, the Internet community devised a plan to preserve addresses by combining class C addresses into supernets. Thus, the organization can have one router on the Internet. The Internet routers only need one entry in their route tables. All traffic to any host on any of the four class C subnets is routed to the single router.

Let's walk through a subnet scenario. Assume that you have decided to use the following as your network ID: 192.168.0.0/16. In a simpler world, you might have considered a situation where your maximum projected number of hosts per subnet had been 8000, and the maximum number of subnets needed was eight. If this were the case, you would subnet this network address into eight networks with 8190 hosts per network, per Table 4-1:

Table 4-1 Simple subnetting

Subnet Number	Subnet Network ID
1	192.168.0.0/19
2	192.168.32.0/19
3	192.168.64.0/19
4	192.168.96.0/19
5	192.168.128.0/19
6	192.168.160.0/19
7	192.168.192.0/19
8	192.168.224.0/19

Now, this is neat and tidy, but not very practical. In fact, you would be hard-pressed to find an organization where equal-sized subnets are the norm. More often, you will find variable-length subnets.

Networks with variable-sized subnets require **variable-length subnet masks (VLSMs)**. Using VLSMs, you can take a single network address and produce subnets of different sizes. These addresses require that the network routing protocols include the subnet mask with the network ID. Unless your routers are using very old protocols, like RIP version 1, this should not be a problem.

In using VLSMs, determine the size of the networks needed and the number of masks needed of each size. For example, let's assume that in your plan you have the following requirements:

- Three subnets needing up to 8000 addresses each
- 30 subnets requiring up to 200 addresses each
- 64 subnets requiring just two addresses each

If you build on the eight networks of the earlier simple scenario, the three networks requiring 8000 addresses could be three subnets from Table 4-1—192.168.128.0/19, 192.168.160.0/19, and 192.168.192.0/19. Then, to create the 30 subnets with up to 254 addresses each, subnet 192.168.224.0/19 by subnetting it further (five bits more). This will actually give you 32 subnets of the correct size. The available network addresses will be 192.168.224.0/24, 192.168.225.0/24, and on up to and through 192.168.255.0/24 (the third octet increments by one). For the extra subnets, you can further subdivide 192.168.225.0/24,

using six additional bits. This range of 64 subnets would have the network addresses 192.168.255.4/30 through 192.168.255.252/30 (the fourth octet increments by four).

The Future: IPv6

The previously discussed 32-bit addressing applies to IP version 4 (IPv4). The next version of IP, version 6 (IPv6), has a 128-bit address, expressed as eight sets of four or fewer hexadecimal digits, with colons as separators between the sets.

An IPv6 address might look like this: 3DF7:AAD:1777:EDCC:91:76B:DEA7:80A3. Notice that there are no leading zeros. Any section that is all zeros can be indicated with a shortened notation. For instance, if the previous example had all zeros in the third set, it could be represented as follows: 3DF7:AAD::EDCC:91:76B:DEA7:80A3. Further, if this same example had all zeros in the fourth set, it would still look like 3DF7:AAD::91:76B:DEA7:80A3. IPv6 calculates the number of contiguous all-zero sets. You cannot have more than one double set of colons in an address.

Don't panic, as complex as these addresses appear, they actually make more sense than the present 32-bit addresses. The new addresses have a wonderful characteristic: They are hierarchical. That means that an IPv6 address actually tells you where in the world a host is. There are three types of IPv6 addresses—unicast, multicast, and anycast—but we limit our discussion to unicast because it is the most important for point-to-point communication.

IPv6 partitions an aggregatable global unicast address into five sections. The leftmost three bits is a format prefix used to identify the type of address. In the example in Table 4-2, the value 001 represents an aggregatable global unicast address. The next 13 bits identify the top-level aggregators (TLAs). These are the Internet public network access points (known as NAPs) that connect the telephone companies with long-distance service providers. The next 32 bits identify the next-level aggregator (NLA), which are assigned addresses by the TLAs. The NLAs are large ISPs such as Pacific Bell Internet. The NLAs, in turn, administer the allocation of addresses for the next level, the site-level aggregator (SLA), which is identified by the next 16 bits in the address. Small ISPs and large organizations, such as universities, are SLAs. The final 64 bits of the address are assigned by the SLAs to their subscribers.

Table 4-2 IP unicast address structure

First 3 bits	Next 13 bits	Next 32 bits	16 bits	64 bits
001	TLA ID	NLA ID	SLA ID	Host Interface ID

Public Addresses When Necessary

If internal hosts need to be addressed directly from the Internet, consider using public addresses. **Public addresses** are addresses assigned to an organization by an ISP or ARIN. Only under special circumstances does an organization need all hosts routable to the Internet. In such a circumstance, one registered TCP/IP address is needed for each host and two for each router. This is generally a less secure design than one using private addresses, although you can use packet filtering on your routers and firewalls to protect the internal network from unauthorized access.

IP Routing in the Intranet Environment

In the intranet environment, an IP router manages traffic between network segments, blocking selected types of traffic and directing traffic to the proper path. A router may be a specialized computer dedicated to routing or a computer running an operating system that includes routing capabilities.

A routing table contains a list of paths to networks. Each line represents a route to a network of which the router is aware. With each network address, the routing table shows its subnet mask and the address to which the router will send all traffic destined for that network. This address will be the address of an adjacent router. Figure 4-7 illustrates a routed network of three subnets connected by two routers, A and B. It also includes simple routing tables for the routers.

It might sound like routing is just a function of routers, but that is not true. Every computer running TCP/IP makes routing decisions. If you open a command prompt on any Windows computer using TCP/IP and type “Route Print,” you will see the local routing table. Figure 4-8 shows a routing table from a Windows 2000 computer.

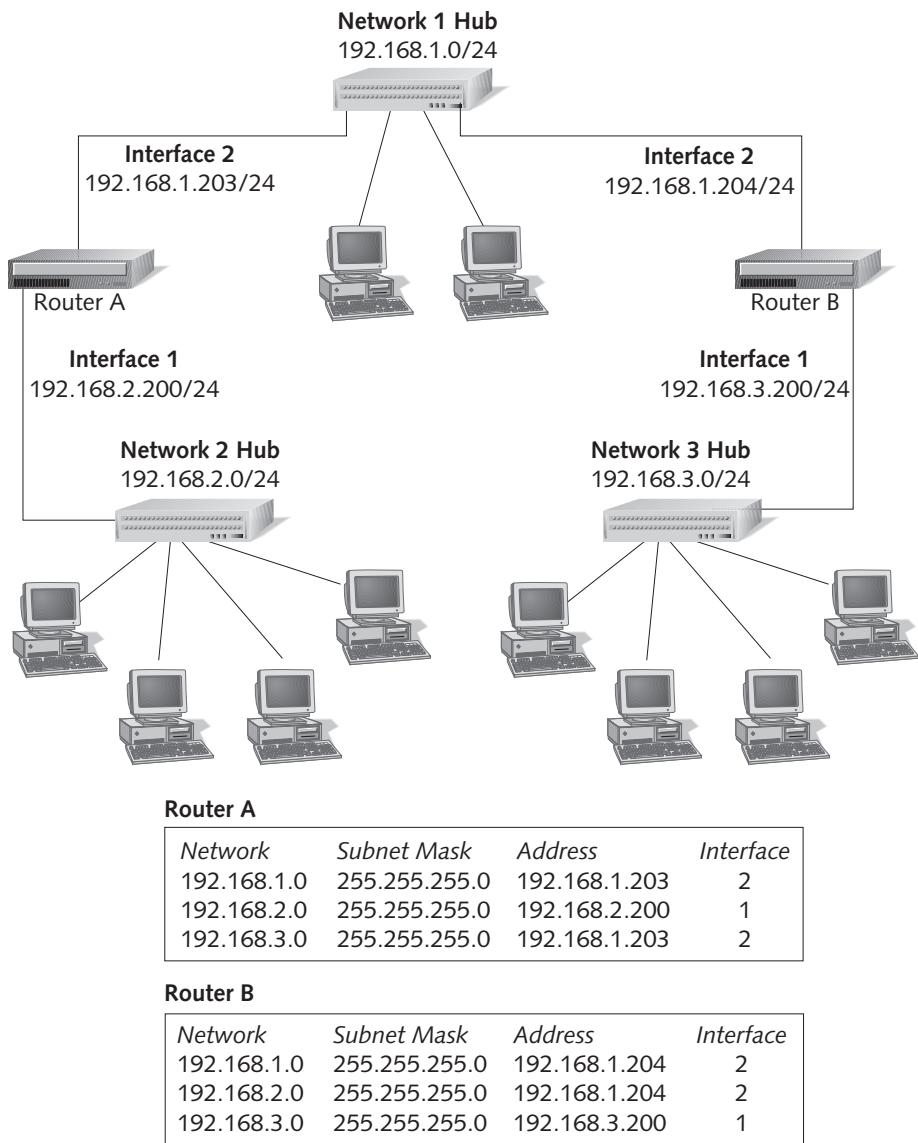
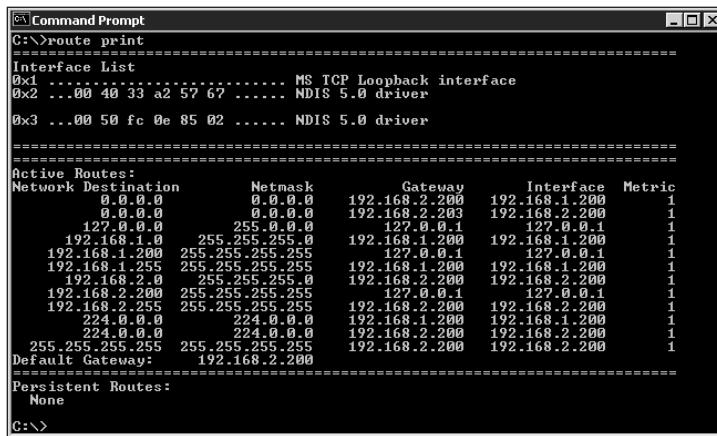


Figure 4-7 Routers and routing tables



```

C:\>route print
=====
Interface List
0x1 ..... MS TCP Loopback interface
0x2 ...00 40 33 a2 57 67 ..... NDIS 5.0 driver
0x3 ...00 50 fc 0e 85 02 ..... NDIS 5.0 driver
=====
Active Routes:
Network Destination      Netmask        Gateway        Interface Metric
          0.0.0.0          0.0.0.0    192.168.2.200  192.168.2.200      1
          0.0.0.0          0.0.0.0    192.168.2.200  192.168.2.200      1
        127.0.0.0        255.0.0.0        127.0.0.1    127.0.0.1      1
        192.168.1.0        255.255.255.0  192.168.1.200  192.168.1.200      1
  192.168.1.200  255.255.255.255        127.0.0.1    127.0.0.1      1
  192.168.1.255  255.255.255.255  192.168.1.200  192.168.1.200      1
        192.168.2.0        255.255.255.0  192.168.2.200  192.168.2.200      1
        192.168.2.200  255.255.255.255        127.0.0.1    127.0.0.1      1
  192.168.2.255  255.255.255.255  192.168.2.200  192.168.2.200      1
        224.0.0.0          255.255.255.255        192.168.2.200  192.168.2.200      1
        224.0.0.0          255.255.255.255        192.168.2.200  192.168.2.200      1
        224.0.0.0          255.255.255.255        192.168.2.200  192.168.2.200      1
        224.0.0.0          255.255.255.255        192.168.2.200  192.168.2.200      1
Default Gateway: 192.168.2.200
=====
Persistent Routes:
  None
C:\>

```

Figure 4-8 Windows 2000 routing table

A **default gateway** is the address of a router on the local network. This is the “gateway of last resort,” meaning that if the host doesn’t have a route to the destination network in its table, it will send the packet to the default gateway. However, it is very common for multiple routers to be on the same network—in which case, a host might have a route to one and the other might be the default gateway. Thus, an IP host does not need knowledge of a default gateway to send packets to remote hosts; it just needs a normal gateway.

Routers use the concept of “longest match,” meaning they always send traffic to the route that has the longest subnet mask. For instance, if you have a route to 10.0.0.0/8 and a route to 10.1.1.0/24 and you want to get to 10.1.1.2, then both routes will get you to your destination, but the /24 is assumed to be closer. A default route is a 0.0.0.0/0. Since you can’t get shorter than /0, a default route is appropriately named the “gateway of last resort.”

Planning and design issues for router protocols are influenced by the size of your network and your IP addressing scheme. The **Routing Information Protocol (RIP)**, although a simple routing protocol to administer, has limits that make it undesirable even for small networks of fewer than 16 subnets. Of the two versions of RIP, 1 and 2, Windows 2000 servers support RIP 2. So, if you plan to use Windows 2000 servers as routers, you will be using the more capable version of RIP, but even RIP 2 is not exactly a great router protocol.

In case you have to work with very old routers, remember that RIP 1 cannot work with VLSMs. Therefore, RIP 1 should only be used if you have just one subnet mask on the network. If such old routers are part of the existing network and you are extending the network and/or planning to use VLSM addressing, plan to replace those routers.

Routing on a multi-homed Windows 2000 computer is accomplished through routing and remote access. Figure 4-9 shows a routing and remote access routing table.

LIVERPOOL - IP Routing Table					
Destination	Network mask	Gateway	Interface	Metric	Protocol
0.0.0.0	0.0.0.0	192.168.2.203	corporate	1	Network management
127.0.0.0	255.0.0.0	127.0.0.1	Loopback	1	Local
127.0.0.1	255.255.255.255	127.0.0.1	Loopback	1	Local
192.168.1.0	255.255.255.0	192.168.1.200	corporate2	1	Local
192.168.1.200	255.255.255.255	127.0.0.1	Loopback	1	Local
192.168.2.0	255.255.255.0	192.168.2.200	corporate	1	Local
192.168.2.200	255.255.255.255	127.0.0.1	Loopback	1	Local
224.0.0.0	240.0.0.0	192.168.2.200	corporate	1	Local
224.0.0.0	240.0.0.0	192.168.1.200	corporate2	1	Local
255.255.255.255	255.255.255.255	192.168.2.200	corporate	1	Local
255.255.255.255	255.255.255.255	192.168.1.200	corporate2	1	Local

Figure 4-9 Routing Table in Routing and Remote Access

Cross-Purposing Your Server

Windows 2000 Server, Advanced Server, and Datacenter Server have moderately sophisticated routing capabilities. Thus, one of these systems can perform routing functions in your routed networks. Although we like the concept of such cross-purposing of servers, some network professionals are adamantly against it for the following reasons:

- Using a server as a router adds unnecessary complexity to an already poor and problematic driver/protocol stack.
- A mix of vendors can complicate things and make troubleshooting and maintenance much more difficult.
- You lose your access to tools such as **Open Shortest Path First (OSPF)** virtual links, complex route redistributions, and timer and metric adjustments.
- You miss out on the literally hundreds of features and options that are available on Nortel and Cisco routers for these protocols.
- Devices that are designed to be routers typically have less overhead and are able to make decisions in hardware, without waiting for other hardware and operating system functions (memory paging to disk, application processes, etc.).
- Servers don't have features such as tunneling, simple port density, WAN interfaces, and ISDN functionality.
- Choosing a server acting as a router or a pure router is not economical and adds complexity to common network problems (such as locked-up print drivers) and maintenance.
- A server acting as a router isn't scalable.

So, if you are planning on creating a cross-purposed server, be aware of these arguments. They may come up in meetings you have with fellow networking professionals.

IP CONFIGURATION STRATEGIES—THE DHCP WAY

If you've gotten this far in the design process, you now have a design for effective use of IP addresses on your network and your design includes meaningful, hierarchical addresses that provide room for growth. Your next step is to decide how each host is going to receive its IP address and other configuration information.

The key labor-saving device in determining IP addresses is the **Dynamic Host Configuration Protocol (DHCP)**, an Internet protocol that allows computers to receive their IP address and configuration over the network from DHCP servers. In the following sections, we examine Windows 2000 DHCP from both the server side and the client side so that you can create an IP configuration strategy that serves your organization well.

DHCP for Windows 2000 and Legacy Operating Systems

DHCP is an Internet standard for automatically assigning addresses to hosts on an IP network. It grew out of an earlier configuration standard, BOOTP, which was designed in 1985. DHCP is enhanced to be a more capable IP configuration tool than BOOTP. For instance, DHCP clients do not have to be listed in a table to receive an address; they also can receive other IP configuration parameters from a DHCP server. However, if you want a host to always be given the same address every time, you may reserve that address by creating a reservation with the DHCP server specifying the client.



BOOTP allows a diskless client machine to discover its own IP address, the address of a server host, and the name of a file to be loaded into memory and executed. The BOOTP administrator creates a table containing clients, their IP addresses, and network configurations. When a BOOTP client boots up on the network, it broadcasts a request for an IP address.

The most commonly used IP configuration parameters that a DHCP server can implement are as follows: Router (Gateway), DNS Servers, DNS Domain Name, WINS/NBNS Servers, and WINS/NBT Node Type.

When you configure a DHCP server, you give it a range of addresses for each subnet that you want it to provide with addresses. Each range of addresses is called a **scope**. The Windows implementation of DHCP includes support for superscopes, which allow you to combine different noncontiguous IP address subnets into a single scope to be applied to the same physical segment. In addition, previously, each scope could have only one set of configuration parameters for all the DHCP clients on that subnet. Now, in the Windows 2000 implementation, DHCP can specify different configurations for clients within the same scope. The configurations can be defined by hardware vendor, by operating system, and by group of users.

Microsoft has defined some vendor-specific options that are currently only supported by NT4 SP3 and Windows 2000 clients. These include:

- **Disable NetBIOS over TCP/IP:** As discussed earlier in this chapter.

- **Release DHCP lease on shutdown:** This will cause a DHCP client to send a release of its DHCP lease when it shuts down.
- **Default Router Metric base:** The client will use this as the base metric for its default gateway(s). A metric is a factor that is used to determine the best network path. A metric might be a **hop count** (number of routers), ticks (a time measurement, such as 1/18 of a second), load, or reliability.

Other features are discussed in the following sections.

Updating DNS for DHCP Clients

The ability to update DNS for DHCP clients is useful only if you are using DNS servers that support dynamic updates of DNS records (Dynamic DNS, or DDNS). If so, your Windows 2000 clients are capable of updating their own DNS records, but legacy clients are not. Figure 4-10 shows a configuration setting of a DHCP server that will accommodate both the new and old clients. Another situation in which you might want to have DHCP update DNS, even for your Windows 2000 clients, is when your DNS servers are configured for secure dynamic updates.

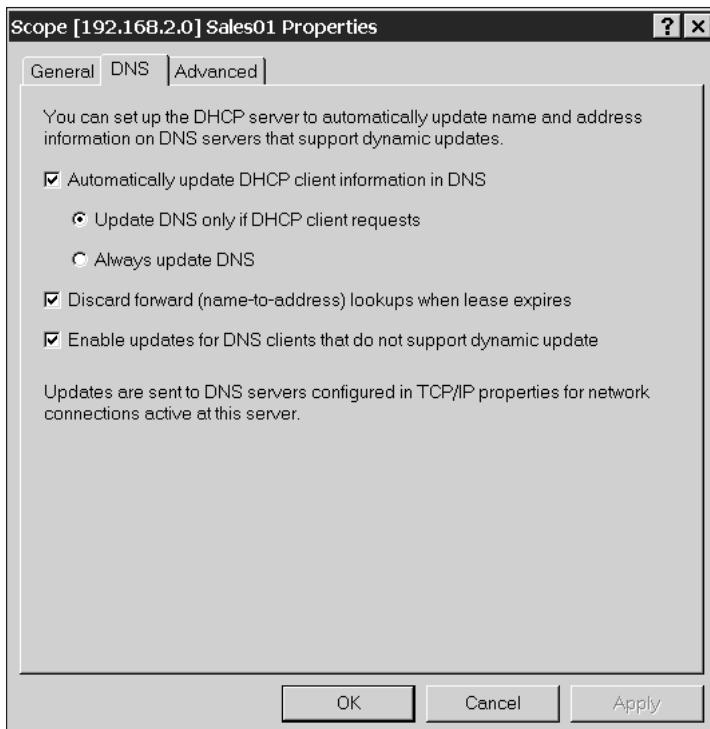


Figure 4-10 DHCP configuration for updating DNS

Multicast IP Address Allocation

Windows 2000 DHCP supports the Multicast Address Dynamic Client Allocation Protocol (MADCAP). MADCAP is the protocol that allows DHCP to assign and configure IP multicast address scopes, which are separate from the standard DHCP scopes. The following are the multicast scope ranges:

239.253.0.0 to 239.253.255.255

239.254.0.0 to 239.254.255.255

239.255.0.0 to 239.255.255.255

DHCP Integration with Remote Access Servers of Remote Users

This feature enables a remote access server to obtain IP address leases from a DHCP server. When a Windows 2000 remote access server initializes, it requests 11 IP addresses from the DHCP server. It will use one of these for its remote access interface and reserve the others to issue to remote access clients. Additional addresses will be requested, as needed, in blocks of 10. Further, if the remote access server has the DHCP relay agent configured with the address of a DHCP server, it will pass on all DHCP configuration information to its clients. Otherwise, the clients only receive the IP address and subnet mask.

Authorization of a DHCP Server by an Active Directory Domain

A problem with DHCP in the past was the danger of “rogue” DHCP servers. New in Windows 2000 DHCP is a feature called “unauthorized server detection.” Although this sounds as if something goes looking for a rogue DHCP server, it is actually up to the server itself to detect that it is not authorized.

When a Windows 2000 DHCP server initializes, it sends out a DHCP Inform packet. It receives replies from other DHCP servers in the form of DHCP Ack messages. The requesting DHCP server compiles a list of all active DHCP servers that respond, along with the root of the Active Directory **forest** used by each server. It then queries the root domain to see if it is authorized, requesting a list of IP addresses of all authorized DHCP servers. If it is on the list, the DHCP service starts. If it is not on the list, it does not start. If it cannot connect, and it does not detect another DHCP server, it will start, but it continues to send DHCP Inform packets every five minutes.

Non-Microsoft DHCP Clients

A Windows 2000 DHCP server supports clients that are compliant with the Internet standards for DHCP. However, test these clients on your test network before changing to a Windows 2000 DHCP server. The client may require something nonstandard that is not supported in Windows 2000, or may not respond well to Microsoft-specific extensions to DHCP. So, the rule is, once again, test it before you use it.

DHCP Clients

You should know that when a DHCP client starts up, it sends out a DHCP Discover packet. This is a broadcast, and any DHCP server hearing this responds. If it does not have an appropriate address for the client, it will respond with a negative acknowledgment (NACK). If it has an appropriate unleased address available, it will send a DHCP Offer packet. When the client receives the Offer packet, it will respond with a DHCP Request packet. If more than one DHCP server made an offer, they now know that their offer was not accepted, and the address they offered is still available. The DHCP server that sent the accepted offer updates its database to indicate that the address has been leased. The server sends a DHCP acknowledgment packet to acknowledge the lease.

DHCP clients receive and hold their DHCP address and configuration for a period of time, called a **lease**, that is configured for each scope on the DHCP server. The lease length is configurable, and the default length in a Windows 2000 DHCP server is eight days. When the client reboots or half the lease time expires, the client sends a Request packet to renew the lease. The server holding the lease responds with a DHCP Ack. This is half the conversation needed to first obtain the lease.

BOOTP Clients

A Windows 2000 DHCP server can support BOOTP clients, but it must be configured to do so. This is actually a scope-level setting, found in the Properties dialog box of each scope on the Advanced tab sheet. Figure 4-11 shows this sheet. Note that the default setting is DHCP only. In addition to the scope-level setting, you will want to go to the properties of the DHCP server and, on the General tab, select Show the BOOTP Table Folder.

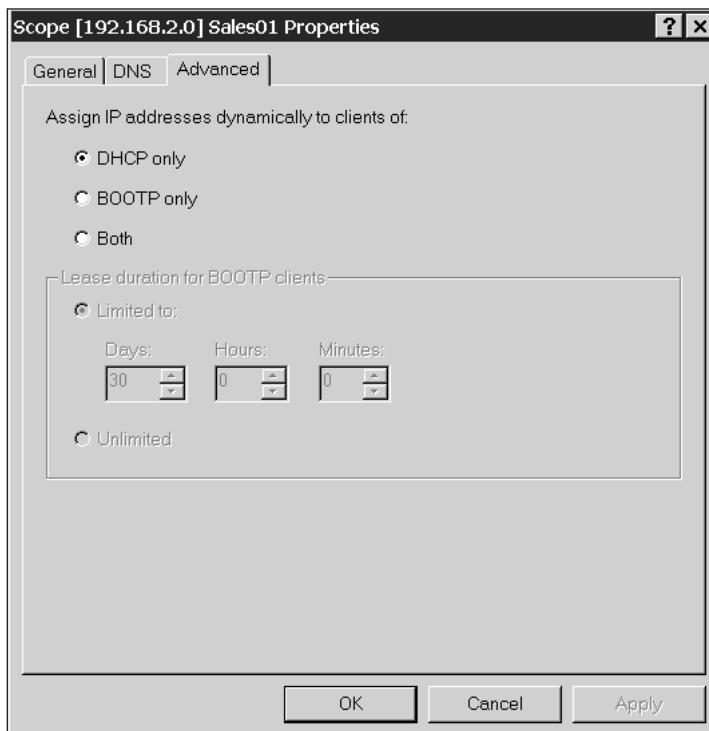


Figure 4-11 DHCP advanced settings

A BOOTP client does not understand the notion of leases, so it will request an address at each startup. Note that the Windows 2000 Remote Installation Service uses BOOTP to communicate with clients to initiate a remote installation.

A Questionable Feature of Windows 2000 DHCP

In Windows 2000 and Windows 98, if a DHCP server cannot be reached, or if lease configuration fails, **Automatic Private IP Addressing (APIPA)** is used. With APIPA, the client will use an address in a special range reserved by Microsoft for use with APIPA. This range is 169.254.0.1 to 169.254.255.254 with a subnet mask of 255.255.0.0. The client will select a number from this range, broadcast it on the subnet to ensure that it is not already in use, and keep it until a DHCP server can be located and a new address leased.

APIPA in a small, single segment network may be fine. If all the resources on that single subnet are using APIPA, they can actually talk to each other. However, in a large multi-segmented network, this one feature could generate a lot of calls to the help desk because these addresses would only be acquired by the clients if the network or DHCP server failed. In such a case, the client would have an address that was useless and that would not allow access to the usual resources on the network, which would most likely

have static addresses. Sounds like fun, doesn't it? That being said, if you run into problems with APIPA, know that you can disable it by using the information contained in Microsoft Knowledge Base article Q244268.

Legacy Windows Operating Systems

Whenever considering how legacy operating systems will work with new features in Windows 2000, remember that the legacy clients' interaction with a Windows 2000 server that has a service with new features will be limited to only what the client is capable of doing. Therefore, legacy DHCP clients will only be able to receive those configuration parameters that their client component understands.

Functionality in a DHCP Design

Once you have arrived at a design for using IP addresses most effectively, you need to determine the following about the addresses needed per subnet:

- The number of addresses that can be dynamically assigned through DHCP
- The number of addresses that can be dynamically assigned if you guarantee the address with a reservation
- The number of addresses that must be static addresses, for servers that run services that require static addresses
- Whether you have non-Microsoft clients or legacy Microsoft clients that will require the DNS server to perform the automatic DNS updates of the client records
- Whether you need multicast scopes

Once you have documented this information for each subnet, you can turn your attention to installing your Windows 2000 DHCP server(s) and configuring scopes per your design. Your design should include the ranges of addresses for use in static configuration, the range to include in your DHCP scopes, the addresses that need to be excluded from distribution, and the hosts that need reservations in DHCP. Figure 4-12 shows the DHCP console with a scope of addresses.

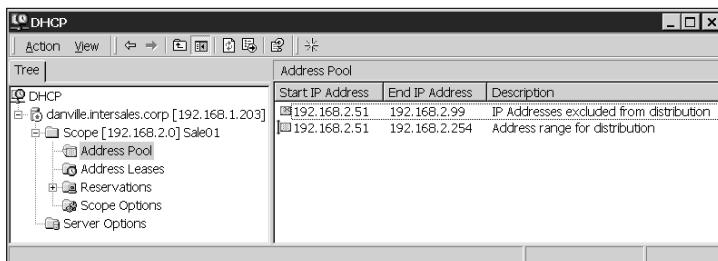


Figure 4-12 DHCP console

Relaying BOOTP and DHCP traffic

DHCP traffic is broadcast traffic, so routers will not forward it. Therefore, a subnet that does not have a DHCP server must have a relay agent on the router or on another computer if DHCP clients on that subnet are to be supported.

A **relay agent** listens for DHCP/BOOTP traffic on a subnet. If it detects a client Discovery or Renewal packet, it will forward the request as a unicast to the DHCP server(s) address(es) entered into its DHCP relay properties. The DHCP sends responses for this client to that relay agent and the agent places the responses on the subnet for the DHCP/BOOTP client to receive.



Do not configure a DHCP relay agent on the DHCP server. Both services use the same ports (UDP ports 68 and 69), and the two services are not compatible on the same machine.

There are several techniques involving server placement that may be used in a routed network. One is to place a DHCP server on each subnet and another is to place the DHCP server(s) on the subnet(s) with the greatest number of hosts. Yet another is to have a DHCP server that is multi-homed, which means it has an interface on two or more networks. We prefer the second option because it would service the local clients; for the remainder of the clients, you could use DHCP relay agents.

Enhancement of a DHCP Design for Availability

To further increase the availability of DHCP, you must use two fault-tolerance techniques for DHCP servers. A fault-tolerance technique is one in which you remove a single point of failure (that single point which, if it fails, will cause problems). One fault-tolerant technique is to use distributed scopes; the other is to use Windows clustering of DHCP servers.

Distributed Scopes

One problem with DHCP servers is that they do not really talk to each other about their leases. In addition, unlike WINS servers, they do not update each other with information from their databases. This makes providing fault tolerance a problem, and it is why distributed scopes are so useful with DHCP.

Here is how one fault-tolerance technique for DHCP works: You place 80% of the addresses for one subnet in a scope on one DHCP server. This is known as the primary DHCP server. Then, you place 20% of the addresses for that same subnet in a scope on another DHCP server, the secondary server. At all times, you are careful not to duplicate any addresses between the scopes.

At this point, if the primary server goes down, clients that need new or renewed leases will get new leases from the surviving DHCP server. If you have planned carefully and

are diligent, you will have the primary DHCP server up and running before enough leases expire to require more than the addresses you have available on the secondary server.

Commonly, for half the subnets, you will make one server primary, the other secondary. For the other half of the subnets, the roles will be reversed. If you set up DHCP servers in a medium-to-large network, you will become pretty good at calculating 80% of the IP addresses. (Hey, this skill might even come in handy on a test!)

Windows Clustering

Windows 2000 supports the use of multiple physical computers, providing a service that appears to be hosted on just one server. DHCP is a cluster-aware service. **Windows clustering** brings the following benefits for the DHCP service and administrator:

- Automatic failover and restart after a failure
- Faster restore of failed servers, because there is only one DHCP database
- No need for distributed scopes

If you decide to use Windows clustering, you must plan for reliable, high-speed connections between the clustered servers.

DHCP for Remote Locations

Placement of DHCP servers on both sides of WAN links is recommended for supporting DHCP in multiple locations. Without such placement, if the WAN link goes down, clients will not receive their IP addresses and will be unable to communicate on the network.

If your remote locations are using dial-up connections and are dialing into Windows 2000 remote access servers, they will receive an IP address and mask for the connection from the server. In addition, if the remote access server has the DHCP relay agent configured with the address of the DHCP server, the clients will receive all other configuration information for their subnet.

Enhancement of a DHCP Design for Security

Security in general has improved in Windows 2000, and for DHCP the best security enhancement has been the detection of unauthorized servers. This only works for Windows 2000 DHCP servers, so if you want this feature, you will need to include Windows 2000 DHCP servers in your design. This is such an important feature for an enterprise that only members of the Enterprise Admins group may authorize a DHCP server. The administrator in the forest root domain is a member of this group.

Speaking of groups, it takes an administrator to install DHCP, an administrator who is a member of the Enterprise Admins group to authorize it in the forest, and membership in two special local groups to perform day-to-day administration and view the configuration of a DHCP server. These groups are DHCP administrators and DHCP users.

DHCP administrators have the right to administer the DHCP server, and DHCP users have read-only access to the DHCP server information. Membership in this second group may be useful for help desk people who may need to check the configuration before creating a ticket to resolve a problem.

Enhancement of a DHCP Design for Performance

Whatever administrators create, they want to enhance. Fortunately, DHCP design is full of opportunities for enhancement. For instance, a multi-homed DHCP server enhances performance because it reduces traffic (since the DHCP server has multiple network adapters, each on a separate subnet). Note that you must use static addresses on all of these NICs.

DHCP can take advantage of multiple processors in the DHCP server, and a faster disk subsystem is probably the best performance enhancement for a DHCP server. Distributed scopes will help performance, especially if the primary DHCP server is placed on the same subnet as the clients, and placing DHCP servers on subnets with the largest number of hosts can be combined with distributed subnets.

When DHCP clients with unexpired leases are off the network for an extended period of time, problems can arise. For instance, if a client leaves without releasing its DHCP address, the address is not available for reallocation to another client until the lease expires or until an administrator manually deletes the lease.

By modifying the lease length, you can control the interval at which a DHCP client renews its lease. Remember that lease renewing creates network traffic, although only a small amount for each client. Modifying the lease also makes addresses available sooner for new DHCP clients—a solution for the problem of clients leaving the network without releasing the lease.

When you are deciding on a lease length, you may find it more art than science, but we can come up with a few rules to help in the process:

- If you have a small network in which the DHCP client computers are rarely moved or changed, you can safely use a very long lease life.
- If there are frequent changes on your network, and clients leaving and returning, you will want a shorter lease length.
- A longer lease length decreases network traffic, but releases addresses later; a shorter lease length causes an increase in network traffic, but releases addresses sooner.

NAME RESOLUTION WITH DNS

Name resolution is extremely important in a network, because it is through name resolution that clients locate network resources without end users having to participate in the process.

DNS is the most important name resolution in a Windows 2000 network. It is required for an Active Directory domain. Anyone who has promoted a Windows 2000 server to a domain controller knows that if DNS is not configured properly, and if the server is not properly configured as a DNS client, a domain controller will not get up and running. The following sections discuss elements that you must consider when creating a functional DNS design.

Pertinent Design Data

The information you have gathered on network and host configuration and distribution will help you to determine the number of DNS servers needed. Each location or, more properly, each Active Directory site should have at least one DNS server. Knowing the number of users will help you to determine the load on the DNS server and perhaps have multiple servers in a site.

Windows 2000 DNS Features

There are three types of zones in Windows 2000 DNS: standard primary, standard secondary, and Active Directory-integrated.

A standard primary zone is a conventional DNS primary zone, meaning that the zone information is stored in a file on the DNS server. A **zone** is a contiguous portion of the Domain Name Space. The DNS server with the primary zone is the only server that may accept changes to the zone. There can be only one primary zone per zone, and a standard primary zone may exist on any Windows 2000 server with the DNS service installed.

A standard secondary zone is a conventional DNS secondary zone, in that the zone information is stored in a file on the DNS server. A DNS server cannot accept changes to the secondary zone except as zone transfers from the DNS server hosting the primary zone. There can be one or more secondary zones for each primary zone, and a standard secondary zone may exist on any Windows 2000 server with the DNS service installed.

An Active Directory-integrated zone can only exist on a domain controller. The zone is stored in Active Directory, so any domain controller in the domain with the DNS service installed can respond to dynamic updates and DNS queries. Each domain controller can accept changes to the zone, therefore making it multi-mastered. An Active Directory-integrated zone is replicated along with Active Directory and is available for dynamic updates from any domain controller in the domain with the DNS service installed.

Windows 2000 Active Directory requires a DNS service properly configured with support for service records (SRV). SRV is the only required new DNS feature. Other new

features in Windows 2000 DNS are not required by Active Directory but are strongly recommended. These include incremental zone updates and dynamic updates of records.

Previously, when there was a change to a zone, the entire zone file was sent in a transfer to secondary zones. Now, only the changed record is sent. This greatly reduces traffic for zone updates. In addition, in the past a DNS administrator used to manually enter all the records in the DNS zones. Now, the DNS service allows clients to update their DNS records, and/or DHCP can update them.



If your network is currently using Unix DNS servers, and it is necessary to keep them in place, check to see if they are running the Berkeley Internet Name Domain (BIND) version of DNS 8.2.1 or later, which supports all these features.

Integration of DNS with DHCP

DHCP can update the DNS records for DHCP clients because of a wonderful new feature implemented in DNS in Windows 2000—dynamic update of DNS. This means the DNS administrator does not have to manually create the DNS records, but can allow them to be created and updated automatically by clients that are capable of this and by DHCP servers for clients that cannot do this.

Additionally, if a DHCP service can be on the same computer as the DNS server hosting the primary zone (or Active Directory-integrated zone), updates can be made to the DNS zone without causing network traffic. Note, however, that this won't always be a viable option. For example, you may be hosting DNS on a Unix server and DHCP on a Windows 2000 server. This configuration would cause network traffic.

Enhancement of a DNS Design for Security

There are several enhancement considerations for a DNS design, two of which we discuss in the following sections.

Secure Dynamic Updates

A DNS zone can be configured to accept only secure dynamic updates. To do this, you must have an Active Directory-integrated zone. Permissions are assigned in Active Directory in the DNS zone container. Give permissions to the computers that need to perform updates, including all statically configured Windows 2000 computers and DHCP servers. Configure the DHCP servers to update both the A and PTR records for clients.

Secure DNS Zone Replication

DNS zones contain computer names and IP addresses. Depending on the security needs of the organization, this information is considered confidential. If your zone replication is occurring over an untrusted network, consider encrypting the traffic.

In the case of standard zone transfers over public networks, encrypt by using IPSec and VPN tunnels implemented through Windows 2000 routing and remote access. Note that Active Directory-integrated zones are, by nature, more secure because they encrypt all replication traffic.

There are some special considerations if your design or existing network includes a screened subnet (DMZ) between your private network and the Internet:

- If you have a need to place a DNS server in the screened subnet, it must not be a domain controller; therefore, you cannot use Active Directory-integrated zones. Also, you should place only secondary zones in the screened subnets.
- If you must place DNS servers in a screened subnet, design your zones so that the servers in the subnet do not expose the entire namespace to queries from the Internet. DNS servers also will perform better if they only have a portion of the namespace.
- Microsoft also recommends configuring firewalls to permit DNS queries only from the Internet, and replication traffic only from the private network; however, you also need to allow for the DNS servers in the screened subnet to forward queries to Internet DNS servers.
- Finally, encrypt zone replication traffic with IPSec or VPN tunnels.

Enhancement of a DNS Design for Availability

The two hardware-based availability enhancements for DNS are (1) using multiple DNS servers for each zone (multiple zones can be on each server) and (2) using Windows server clusters. Note that the latter does not help you with your remote locations, because of the need for persistent high-speed connections between the servers in a cluster.

In addition, there are certain “tweaks” you can use to enhance your design:

- In an Active Directory-integrated zone, adjusting the Active Directory replication schedule between sites will make changes available sooner throughout the domain.
- For standard DNS zones, having multiple secondary zones for each primary zone increases availability.
- The use of incremental zone transfers (IXFR) improves availability by not sending out the entire zone, only the changes. Therefore, it makes the changes available sooner.

Enhancement of a DNS Design for Performance

If you have Windows DNS servers hosting standard primary zones and non-Windows 2000 servers hosting secondary zones, the secondary zones may require complete zone transfers (AXFR). This will hurt performance. To mitigate performance degradation, consider the following:

- To reduce query resolution time, install caching-only servers at remote locations connected by low-speed WAN links, especially if the available bandwidth cannot handle the addition of DNS zone replication traffic.
- Caching-only servers also work well when the DNS zone information does not change frequently. At installation, a DNS server is a caching-only server.
- To minimize the size of the zone databases, delegate portions of the namespace to additional DNS primary servers.



There is a capability of DNS that we will be recommending later in the book for load balancing of other services. This is Round Robin DNS, which has been part of DNS for some time and is a trusted tool of network administrators. When you employ Round Robin DNS, you create records in zones for several servers running the same service. The servers have unique names and unique IP addresses, but in DNS, they are aliased to the identical server name. When DNS sees the same server name with multiple IP addresses, it will resolve successive queries by moving through the list, thus providing a form of load balancing.

NAME RESOLUTION WITH WINS

When a user initiates a request to access a network resource, the destination computer is probably referenced by a special name. In a TCP/IP network, this is most likely to be a DNS name, but could alternately be a NetBIOS name. Before a packet can be sent to a network destination, this name must be resolved to a logical address, which, in the case of TCP/IP, will get the packet to the appropriate subnet.

NetBIOS is a protocol developed by IBM in the 1970's. When Microsoft adopted NetBIOS for its network, they had a close relationship with IBM. Windows 2000 uses Internet-style domain names as its namespace for Active Directory; therefore, DNS is the primary name resolution method for Windows 2000 clients in an Active Directory domain. However, as long as we have non-Windows 2000 clients on our networks, and old applications that require NetBIOS, we will have to include NetBIOS name resolution in our network designs.

The Functional WINS Design

To create a functional WINS design, you must consider the features in the Windows 2000 implementation of WINS, the number of WINS servers needed, WINS replication partnerships, the impact of WINS traffic on slow WAN links, the integration of WINS with other services, and pertinent design data.

The information you have gathered previously in this design process on network and host configuration and distribution will help you to determine the number of WINS servers needed. Each location with non-Windows 2000 clients and/or NetBIOS applications should have at least one WINS server. Knowing the number and distribution of

users will help you to determine the load on the WINS server and perhaps indicate the need for multiple servers in a site.

WINS Functionality and Features

WINS is a NetBIOS Name Server (NBNS). You use WINS for NetBIOS support in a routed network, because without WINS, you will have to use LMHOSTS files, which are unwieldy in a large network. The WINS server provides many functions. We describe each in turn in the following sections.

Name Registration

Each WINS client is configured by providing it with the WINS server address as a parameter of the TCP/IP properties for that computer. This can be done manually or through DHCP. Windows 2000 has increased the number of WINS server addresses that can be provided from two to 12. As a WINS client starts up, it registers its name, IP address, and all its network services with a WINS server.

Name Resolution

The node type controls the default client behavior during NetBIOS name resolution. One node type, Hybrid, is the default node type for a statically configured Windows WINS client. When a Hybrid node client needs to perform name resolution, it attempts the following steps, in order, until the name is resolved or until all methods fail:

1. The client checks the NetBIOS name cache to see if the name has been resolved recently.
2. The client sends a name resolution query to the first WINS server on its list of WINS servers. If that server does not respond, it sends a query to the next server on its list. Once a WINS server responds, no other WINS server is contacted, even if the server that responds cannot resolve the query.
3. If the name is still not resolved, it sends a NetBIOS broadcast to request the name resolution.
4. It checks the LMHOSTS file, if it exists.
5. It checks the HOSTS, if it exists.
6. It queries DNS.

Since the latter methods are not likely to produce positive results, if the name cannot be resolved by WINS or a NetBIOS broadcast, it can take several seconds to get the negative response.

Registration Renewal

When a WINS client registers with WINS, the server sends the client a packet with a Time to Live (TTL) value, which indicates when the client records will expire and need

to be renewed. If the registration is not renewed by the end of the TTL, it expires in the WINS database, and will eventually be removed. The default TTL value is six days. Whenever multiple WINS servers are used, you should use the same TTL on all of the servers that are replication partners.

Name Release

When a WINS client shuts down (properly) or when a user runs the “nbtstat” command with the new “-RR” parameter, the WINS client sends a name release message to its WINS server. The WINS server automatically marks the entry for that client as released. A record that is released can be replaced by a registration request from a client with the same name but a different IP address. If the record remains released for a certain interval, the server **tombstones** it. Thus, the record is marked for eventual deletion, and the server modifies the version ID and notifies its replication partners of the change.

Replication Partners

If you have multiple WINS servers, you may configure them to replicate their databases with each other. This provides a consistent database of all registered NetBIOS resources on the network. Each server must be configured to replicate with another server as a push, pull, or push/pull partner. The latter is the default, and is the recommended configuration. A **pull partner** requests data to be sent to it from other WINS servers. A **push partner** sends data to other WINS servers based on the number of changes to the database. A **push/pull partner** combines the two; it requests changes from partners at an interval and pushes changes to partners when there are changes to the database.

Burst-Mode Name Registration

There are times when an unusual number of WINS registrations may occur (a burst of activity). One such example is when there is a power failure. When the power comes back on, the WINS clients will all attempt to register at once. The WINS server can get overwhelmed at this time and may fail to acknowledge registrations or fail to properly enter the records in the database. If it does not acknowledge registrations, WINS clients continue to attempt to register, compounding the problem. If it fails to properly enter the records in the database, name resolutions will fail.

WINS has a new technique for handling this—burst-mode name registration. Essentially, it tells each client to go away and come back when the server is not so busy. It does this by issuing an acknowledgment to each client with a shorter than usual TTL. This makes the client happy, thinking that everything is OK until the end of the TTL, when the client will re-register. At this time, the record gets properly registered.

Persistent Connections and Manual Tombstoning

WINS servers can now be configured to maintain persistent connections with replication partners, which increases reliability and performance of WINS replication.

Additionally you can now manually mark a WINS record as tombstoned. This change, as with all changes, will be replicated to a WINS replication partner.

Enhanced Filtering and Record Searching

With several entries for each registered client, the WINS database can grow to be huge. The former administrative interface was very awkward to use, showing you much more information than you needed. The new WINS console makes it much easier to locate records in the WINS database, only showing records that fit the criteria you specify. Figure 4-13 shows a WINS console with host records displayed in the contents pane.

For a test network or other small network, when you really do want to see it all, the trick to seeing all the records in the database at one time is to open the WINS console, right-click the Active Registrations folder, and then click Find by Owner. In the Find by Owner dialog box, select All Owners, and then click the Find Now button.

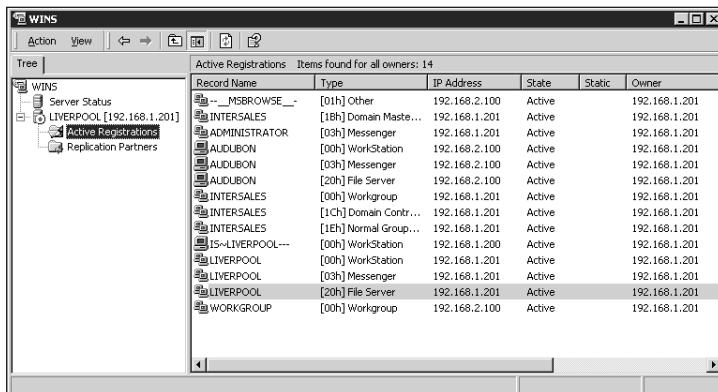


Figure 4-13 WINS console

An owner is the WINS server that registered the record. So, if you only have one WINS server, or are interested in viewing the records owned by a particular server, select This Owner in the previous step.

Integration of WINS with DHCP and DNS

You should configure your DHCP scopes to provide DHCP clients with WINS server addresses. This configuration does not have any negative consequences, only benefits, because a WINS server will handle the resulting NetBIOS registrations and queries in a speedy manner.

Integrating WINS with DNS can be more problematic. Consider the situation in which a DNS server is configured as a WINS client and you have configured a zone on that DNS server to use WINS forward lookup. In such a case, when the DNS server is unable to resolve a DNS query, it will query WINS for the resolution of the server name portion of

a DNS name. Unfortunately, this query can negatively affect performance, especially if the WINS lookup fails.

Enhancement of a WINS Design for Security

WINS databases contain computer names and IP addresses. Depending on the security needs of the organization, this information is considered confidential. If your WINS replication is occurring over an untrusted network, consider encrypting the traffic by using IPSec and VPN tunnels implemented through Windows 2000 Routing and Remote Access.

Enhancement of a WINS Design for Availability

The first enhancement option is to have multiple WINS servers, preferably one or more at each site. You can also place WINS on Windows 2000 server clusters. Note, however, that server clusters do not help you with your remote locations, because of the need for persistent high-speed connections between the servers in a cluster. The second enhancement option is to modify replication parameters. In addition, you can reduce the length of time between WINS database replications by providing server hardware that supports WINS.

Enhancement of Response Time to Requests

Although this is an area that has been improved in the WINS service, you can still improve the responsiveness of a WINS server to client requests by improving the hardware with multiple processors, additional memory, a high-performance disk subsystem, a high-performance network interface card, multiple WINS servers, and WINS servers on either side of WAN links. In addition, you can improve response time by enabling burst-mode name registration on the server if extremely high demand is predicted and by distributing clients across multiple servers for load balancing.

Enhancement of WINS Replication

When WINS replication must occur across WAN links, it can affect the available bandwidth for all traffic. However, configuring schedules for the replication traffic across WAN links may affect convergence. **Convergence** is the point at which the databases of all replication partners match. Although convergence is desirable for availability, you may have to balance the need to reach convergence with the requirement for performance.

To balance convergence with performance, you can control the time and frequency of replications through WINS replication schedules so that they occur during off-peak hours or at an interval that does not compete with user traffic. You also can consider using persistent connections to maintain connections between replication partners.

CHAPTER SUMMARY

The first part of this chapter focused on TCP/IP security features and performance enhancements. We then moved to TCP/IP addressing strategies, which is at the root of everything in this book. We gave you information and tools to help you analyze both present and future subnet requirements. We also gave you knowledge for designing an IP addressing and implementation plan based on the number of hosts and subnets available as well as the need for public or private network access. You also learned to integrate Windows 2000 IP routing capabilities into existing networks.

The second section focused on DHCP configuration strategies for both Windows 2000 and legacy operating systems. You studied DHCP in a routed environment and for remote locations. Of course, you learned more about enhancing DHCP design for functionality, security, availability, and performance.

The next section focused on DNS, the DHCP/DNS synergies, and how to enhance a DNS design for the four basic parameters of a network. The last section focused on WINS, the WINS/DNS synergies, and how to enhance a WINS design for functionality, security, availability, and performance.

KEY TERMS

Application-layer packet filtering — Allows filtering of packets on a host-by-host basis.

Automatic Private IP Addressing (APIPA) — A feature of the Microsoft TCP/IP stack since Windows 98. With APIPA, when a client configured to receive its address automatically does not receive a response from a DHCP server, it will use an address in a special range reserved by Microsoft for use with APIPA. This range is 169.254.0.1 to 169.254.255.254 with a subnet mask of 255.255.0.0. The client will select a number from this range, broadcast it on the subnet to ensure that it is not already in use, and keep it until a DHCP server can be located and a new address leased.

bastion hosts — A gateway between an inside network and an outside network designed to defend against attacks aimed at the inside network.

convergence — The point at which the databases of all replication partners match.

default gateway — The address of a router on a local network.

demilitarized zone (DMZ) — A screened subnet between firewalls.

Dynamic Host Configuration Protocol (DHCP) — An Internet protocol that allows computers to receive their IP address and configuration over the network from DHCP servers.

failover — In a server cluster, a method by which a server automatically takes over for a failed server.

forest — A forest consists of one or more trees, each of which contains one or more domains that share the same schema, configuration, and global catalog.

hop count — A metric used in routing that indicates the number of routers that a packet must traverse in a certain route.

hosts — A name commonly used to refer to computers in a TCP/IP network.

Internet Control Message Protocol (ICMP) router discovery — Allows a host to discover a router automatically, in spite of not having a default gateway configured in its TCP/IP properties.

Internet Security Association Key Management Protocol (ISAKMP) — An IPSec protocol that provides the method by which two computers can agree on a common set of security settings. It also provides a secure way for them to exchange a set of encryption keys to use for their communication.

IP Security (IPSec) — A set of standards developed by the IETF for the next version of IP—IPv6—and as an optional extension to IPv4. It is included in Windows 2000. IPSec allows for authentication of the source and destination hosts before data is sent. It also allows for the encryption of the data packets during transmission.

jitter — The period frequency displacement of a signal from its ideal location.

lease — The period of time for which DHCP clients receive and hold their DHCP address and configuration information.

Oakley — A key determination protocol of IPSec that uses the Diffie-Hellman key exchange algorithm.

Open Shortest Path First (OSPF) — A routing protocol support by Windows 2000 that is preferred over RIP for larger networks. OSPF works best in a hierarchically designed network.

port — An identifier used in a TCP/IP packet to determine the program or service that is sending or receiving data. Ports are associated with protocols, such as TCP or UDP. For instance, TCP port 20 identifies File Transfer Protocol (FTP) data.

public addresses — Addresses assigned to an organization by an ISP or ARIN.

pull partner — Requests data to be sent to it from other WINS servers.

push partner — Sends data to other WINS servers based on the number of changes to the database.

push/pull partner — Requests changes from partners at an interval and pushes changes to partners when there are changes to the database.

Quality of Service (QoS) — A name for a set of components by which Windows 2000 provides bandwidth reservation capability.

relay agent — A computer that listens for DHCP/BOOTP traffic on a subnet.

Routing Information Protocol (RIP) — A simple routing protocol for small internetworks of less than 16 subnets. Windows 2000 supports RIP version 2 for IP and IPX protocols, but has a limit of 15 subnets.

scope — A contiguous range of addresses for a single subnet.

security association (SA) — The combination of the security method agreed upon and the keys the method uses.

subnetting — The act of taking a network address, such as 192.168.0.0/16, and borrowing bits from the host portion to subdivide this single network address into multiple network addresses.

supernetting — Borrowing bits from the network portion to combine several network addresses into one.

tombstoning — Marking something in a database to eventually be deleted.

transport mode — The mode in which IPSec can be used to authenticate and/or encrypt communications between computers without using a tunnel.

tunnel mode — The mode in which IPSec will encapsulate IP packets and optionally encrypt them.

variable-length subnet masks (VLSMs) — Used to produce subnets of different size from a single network address.

Windows clustering — The use of multiple physical computers to provide a service that appears to be hosted on just one server.

zone — A contiguous portion of the Domain Name Space.

REVIEW QUESTIONS

1. You are an outside consultant designing an IP addressing strategy for your new customer's company, which does not have internal IT personnel. They have 300 network hosts. They expect 30% growth in the number of hosts over the next 18 months. You have recommended private IP addresses, but their operations manager has questioned why they should do this, since they have a class C network address. Which of the following justify your recommendation? (Choose all that apply.)
 - a. Private addresses route on the Internet.
 - b. Use of private addresses provides better security.
 - c. Private addresses are more secure because they do not route on the Internet.
 - d. Private addresses give you more flexibility as the organization grows internally.
 - e. A single class C will not give them enough host addresses.
2. In the scenario in Question 1, what would be an indicator that you should not recommend private IP addresses?
3. You have been selected to head a team designing an IP addressing scheme for your company's 10,000-host intranet. What are some of the criteria you will use for this design? (Choose all that apply.)
 - a. random
 - b. meaningful
 - c. hierarchical
 - d. alphabetical

The next five questions are based on the following case study:

The ABC Corporation has hired your consulting company to aid in the restructuring of its existing TCP/IP network. They have four sites. The headquarters is in Philadelphia, with a regional office in Boston and branch offices in Baltimore and Atlanta. The number of hosts per site is as follows:

Philadelphia: 1209

Boston: 735

Baltimore: 589

Atlanta: 150

There is a 168 Kbps fractional T-1 connection from Atlanta to the Baltimore network, and a T-1 connection from Baltimore to Philadelphia. There is also a T-1 connection from Boston to Philadelphia. All locations access the Internet through the proxy server/firewall in Philadelphia. This server has a leased IP address.

They are using private addresses on the internal network, and the routers at the four locations currently have the following addresses on the private interfaces: 172.20.32.1/19, 172.20.64.1/19, 172.20.96.1/19, 172.20.128.1/19

The company is replacing the existing routers and installing new routers. At some sites, they may have to add routers and further subnet the network, because the new routers will only have 220 hosts per subnet.

4. How many public addresses are required for the current network?
5. How many subnet addresses are needed for the private network?
6. Which network prefix/subnet mask would work best in your design?
 - a. /19
 - b. /15
 - c. /24
 - d. /30
7. What is the total number of host addresses needed?
8. How many hosts per subnet does the network prefix/subnet mask allow?
9. A Windows 2000 DHCP server cannot provide IP leases for legacy clients. True or False? Explain your answer.
10. You are designing a network for a routed environment all on one campus, with a great deal of network traffic. There are 26 subnets with a total of 3081 client hosts and 30 servers. All desktop computers will be DHCP dynamic clients. All servers are being upgraded to Windows 2000. Recently, the file and print traffic has been practically eliminated across subnets by placing file and print servers in each subnet. These servers will all be replaced with new hardware and Windows 2000 Server or Advanced Server.

The four remaining servers are three NT domain controllers and one Exchange 5.5 Server on a single subnet. The domain controllers will also be upgraded to Windows 2000 domain controllers. All servers have static IP addresses.

Currently, greater than 80% of the file and print traffic is within the subnets. Of these subnets, five have 180 to 210 clients. Twenty-one subnets have 130 to 165 hosts. The routers will not be configured to route DHCP/BOOTP traffic. You are making design decisions concerning DHCP.

You have decided to use several DHCP servers, adding six machines for this function, and to provide additional file and print servers. Based on the information provided, where can you place the servers to provide availability?

11. In the previous scenario, how can you provide for fault tolerance?
12. What measures can you take in the scenario in Question 10 to provide security?
13. What steps can you take to enhance a DHCP design for performance?
14. Add this information to the information in Question 10: The company has merged with another company and is adding four regional offices in four different cities. The regional offices have from 200 to 400 users at each site. They will place a single DHCP server at each of these sites, autonomously providing the IP configuration for the DHCP hosts at the sites. There will be Windows 2000 file and print servers on each subnet. Each site will have a new domain controller in the corporate Windows 2000 domain. Describe a functional DNS design for their expanded network.
15. How can you enhance a DNS design for security?
16. How can you enhance a DNS design for availability?
17. How can you enhance a DNS design for performance?
18. To the scenario from Question 10 and Question 14, add this information: The company has several legacy applications that depend on NetBIOS to request access to network resources. It will also take about a year before all the desktop computers are converted from Windows 9x and NT to Windows 2000. Meanwhile, clients needing WINS support are distributed on all subnets. What do you suggest as a WINS strategy that will provide both functionality and security?
19. Building on the scenario in Question 18, how would you enhance your WINS design for availability and performance?
20. Explain how to integrate WINS, DHCP, and DNS.

HANDS-ON PROJECTS



Project 4-1 Examining a Routing Table Using the Route Print Command



For this project, you will need a computer running Windows 2000 Server.

When you are troubleshooting network connection problems, one of the command line utilities you will use is the route command. In this project, you will use the route command with the print option to view the route table of the Windows 2000 computer you are using.

1. If your server is not powered up, power it up now.
2. Press **Control/Alt/Delete** to display the Security dialog box titled Log on to Windows.
3. In the User name box, type **administrator**.
4. In the Password box, type **password**. (If this does not work, ask the instructor for the password.)
5. In the Log on to box, use the selection arrow to select **INTERSALES**. (This, too, will depend on the classroom configuration.)
6. Press **Return**.
7. When the desktop appears, click the **Start** button on the taskbar.
8. Click **Run**.
9. In the Open box of the Run dialog box, type **cmd**.
10. Click the **OK** button.
11. At the command prompt, type **route print**, and then press **Enter**. To interpret the information you see on your screen, you will use the online help in Windows 2000 Server.
12. Click the **Start** button on the taskbar.
13. Click **Help**.
14. In the Windows 2000 window, locate the text box titled “Type in the word(s) to search for” on the Search tab.
15. In the box, type (including the quotes) “**route print**”.
16. Click the **List Topics** button.
17. When the search is completed and the results appear, locate the Select topic box and double-click **The Windows 2000 IP Routing Table**.

18. Use the explanation and table in this article to interpret the route table you have on your desktop.
19. When you have finished, close all windows.



Project 4-2 Installing, Authorizing, and Configuring a DHCP Service

4



For this project, you will need a computer that is running Windows 2000 Advanced Server and is a member server in the "intersales.corp" domain. It must have a static IP address and use the server Liverpool for DNS name resolution. You will also need to know the location of the Windows 2000 Server source files. Your instructor will give you this information, which you will need in Step 14.

You are setting up a Windows 2000 server to be a DHCP server. You need to install the DHCP service, authorize the DHCP server, and then create a scope of addresses.

1. If your server is not powered up, power it up now.
2. Press **Control/Alt/Delete**.
3. In the User name box, type **administrator**.
4. In the Password box, type **password**.
5. In the Log on to box, use the selection arrow to select **INTERSALES**.
6. Press **Return**.
7. When the desktop appears, click the **Start** button.
8. Point to **Settings**, click **Control Panel**, and then double-click **Add/Remove Programs**.
9. In the Add/Remove Programs window, click **Add/Remove Windows Components**.
10. In the Components: section of the Windows Components Wizard, use the scroll bar to scroll down the component list until the words "Networking Services" appear.
11. Click the words **Networking Services** (do not click the check box), and then click the **Details** button on the right, below the Components box.
12. In the Networking Services window, click the check box for **Dynamic Host Configuration Protocol (DHCP)**, and then click **OK**.
13. When you are returned to the Windows Components Wizard, click **Next**.
14. If the Insert Disk dialog box appears, click **OK**, enter the path provided by your instructor, and then click **OK**.
15. At the completion of the installation, click **Finish**.
16. Click the **Close** button.
17. Close all open windows.

In the following steps, you will authorize your DHCP server in the “intersales.corp” domain. You begin by opening the DHCP console from the Administrative Tools menu.

1. Click the **Start** button on the taskbar, point to **Programs**, point to **Administrative Tools**, and then click **DHCP**. The DHCP console opens on the desktop.
2. Right-click **DHCP**.
3. Click **Manage Authorized Servers**.
4. In the Manage Authorized Servers dialog box, click the **Authorize** button.
5. In the Authorize DHCP Server dialog box, type the name or IP address of your server.
6. Click the **OK** button.
7. In the DHCP message box, click the **Yes** button to confirm authorization.
8. In the Manage Authorized Servers dialog box, click the **Close** button.
9. In the DHCP console, click your server name. A red arrow should appear on the server icon to the left of your server name. From this point, you will create and configure a DHCP scope.
10. Right-click your server name in the DHCP console. A context menu should appear.
11. Click **New Scope**.
12. In the New Scope Wizard, click the **Next** button.
13. On the Scope Name page, type your name in the Name box.
14. Click the **Next** button.
15. In the IP Address Range page of the New Scope Wizard, type the following address in the Start IP address box: **10.x.0.51** (where *x* is your student number).
16. In the IP Address Range page of the New Scope Wizard, type the following address in the End IP address box: **10.x.0.254** (where *x* is your student number).
17. In the IP Address Range page of the New Scope Wizard, change the value in the Length box to **16**. Notice that the value in the Subnet Mask box changes to **255.255.0.0**.
18. Click the **Next** button.
19. In the Add Exclusions page of the New Scope Wizard, type **10.x.0.51** in the Start IP address box (where *x* is your student number).
20. In the Add Exclusions page of the New Scope Wizard, type **10.x.0.99** in the End IP address box (where *x* is your student number).
21. Click the **Add** button. The range of addresses appears in the Excluded address range box.

22. Click the **Next** button. The Lease Duration page of the New Scope Wizard appears. Notice that the default lease duration is eight days. You would change this if you desire a different lease duration.
23. Click the **Next** button. The Configure DHCP Options page of the New Scope Wizard appears. We will configure options.
24. Be sure that the **Yes, I want to configure these options now** option is selected.
25. Click the **Next** button.
26. The Router (Default Gateway) page of the New Scope Wizard appears. For this exercise, we will assume that the router on your network is 10.x.0.1 (where x is your student number).
27. In the IP address: box, type **10.x.0.1** (where x is your student number).
28. Click the **Add** button. Notice that the address moves to the bottom box.
29. Click the **Next** button. The Domain Name and DNS Servers page of the New Scope Wizard appears.
30. In the Parent domain: box, type **intersales.corp**.
31. In the Server name box, type **Liverpool**.
32. In the IP Address box, type the address of Liverpool.
33. Click the **Add** button. This moves the address into the box below for DNS servers.
34. Click the **Next** button. The WINS Servers page of the New Scope Wizard appears.
35. In the Server name box, type **Liverpool**.
36. In the IP Address box, type the address of Liverpool.
37. Click the **Add** button. This moves the address into the box below for WINS servers.
38. Click the **Next** button. The Activate Scope page of the New Scope Wizard appears.
39. Select the **Yes, I want to activate this scope now** option.
40. Click the **Next** button. The Completing the New Scope Wizard page of the New Scope Wizard appears.
41. Click the **Finish** button. The DHCP console appears. A green arrow should appear on the server icon to the left of your screen name.
42. Click **Scope [10.x.0.0]**, where x is your student number.
43. Click **Address Pool** to see the range of addresses in the scope and the addresses excluded.
44. Click **Scope Options** to see the scope options you configured.
45. Close the DHCP console window.



Project 4-3 Installing and Configuring the DNS Service



For this exercise, you will need a computer running Windows 2000 Advanced Server that is a member server in the “intersales.corp” domain. It must have a static IP address and use the server Liverpool for DNS name resolution. You will also need to know the location of the Windows 2000 Server source files. Your instructor will give you this information, which you will need in Step 13. The Liverpool server should have a forward DNS lookup zone for the intersales.corp domain, and it should have a reverse lookup zone for subnet 192.168.1.0/24 that is configured to allow transfers to any server. If this information is different for your lab, your instructor will provide you with the correct network address for Step 27.

You are setting up a Windows 2000 server to be a DNS server. You will install the DNS service, and then create a standard secondary zone.

1. If your server is not powered up, power it up now.
2. Press **Control/Alt/Delete**.
3. In the User name box, type **administrator**.
4. In the Password box, type **password**.
5. In the Log on to box, use the selection arrow to select **INTERSALES**.
6. Press **Return**.
7. When the desktop appears, click the **Start** button.
8. Point to **Settings**, click **Control Panel**, and then double-click **Add/Remove Programs**.
9. In the Add/Remove Programs window, click **Add/Remove Windows Components**.
10. In the Components: section of the Windows Components Wizard, use the scroll bar to locate and click the words **Networking Services**, and then click the **Details** button.
11. In the Networking Services window, click the check box for Domain Name System (DNS), and then click **OK**.
12. When you are returned to the Windows Components Wizard, click **Next**.
13. If the Insert Disk dialog box appears, click **OK**, enter the path provided by your instructor, and then click **OK**.
14. At the completion of the installation, click **Finish**.
15. Click the **Close** button, and then close all open windows. The DNS service is installed. You do not have to restart the machine. You are ready to configure the DNS service.

16. Click the **Start** button on the taskbar, point to **Programs**, point to **Administrative Tools**, and then click **DNS**.
17. In the DNS console, right-click your **server name** in the left pane of the window.
18. Click **Configure the server**.
19. In the Welcome to the Configure DNS Server Wizard page of the Configure the DNS Server Wizard, click the **Next** button.
20. In the Forward Lookup Zone page of the Wizard, click **Yes, create a forward lookup zone**, and then click the **Next** button.
21. In the Zone Type page of the Wizard, click **Standard secondary**, and then click the **Next** button.
22. In the Zone Name page of the Wizard, type **intersales.corp**, and then click **Next**.
23. In the Master DNS Servers page of the Wizard, type the IP address of Liverpool in the IP address box. Depending on the classroom setup, this should be **192.168.1.200**.
24. Click the **Add** button, and then click the **Next** button.
25. In the Reverse Lookup Zone page of the Wizard, click **Yes, create a reverse lookup zone**, and then click the **Next** button.
26. In the Zone Type page of the Wizard, select **Standard secondary**, and then click the **Next** button.
27. In the Reverse Lookup Zone page of the Wizard, confirm that **Network ID** is selected and enter the following number (or a network address provided by your instructor) into the Network ID box: **192.168.1**.
28. Click the **Next** button.
29. In the Master DNS Servers page of the Wizard, type the IP address of Liverpool in the IP address box. Depending on the classroom setup, this should be **192.168.1.200**.
30. Click the **Add** button, click the **Next** button, and then click **Finish**.
31. In the DNS console, expand your server object in the tree pane, and then expand **Forward Lookup Zones** to expand the folders below it.
32. Intersales.corp should appear in a folder below Forward Lookup Zones. If it does not, right-click **Forward Lookup Zones**, and then click **Refresh**.
33. Double-click **intersales.corp** and view the records in the right pane of the console window. You should see the servers that are registered and the service records (folders with names that begin with an underscore).
34. Close the DNS console.



Project 4-4 Installing and Configuring a WINS Service



You will need a computer running Windows 2000 Advanced Server that is a member server in the “intersales.corp” domain. It must have a static IP address and use the server Liverpool for DNS name resolution.

You are installing the WINS service on a Windows 2000 server to support legacy clients and applications that still use NetBIOS. You need to install the WINS service, and then open the WINS console and verify that the WINS service is running on the server.

1. If your server is not powered up, power it up now.
2. Press **Control/Alt/Delete**.
3. In the User name box, type **administrator**.
4. In the Password box, type **password**.
5. In the Log on to box, use the selection arrow to select **INTERSALES**.
6. Press **Return**.
7. When the desktop appears, click the **Start** button on the taskbar.
8. Point to **Settings**, and then click **Control Panel**.
9. In Control Panel, double-click **Add/Remove Programs**.
10. In the Add/Remove Programs window, click **Add/Remove Windows Components**.
11. In the Components: section of the Windows Components wizard, use the scroll bar to scroll down the component list until the words “Networking Services” appear.
12. Click **Networking Services**, and then click the **Details** button on the right, below the Components box.
13. In the Networking Services window, click the check box for Windows Internet Name Service (WINS), and then click **OK**.
14. When you are returned to the Windows Components Wizard, click **Next**.
15. If the Insert Disk dialog box appears, click **OK**, enter the path provided by your instructor, and then click **OK**.
16. At the completion of the installation, click **Finish**.
17. Click the **Close** button on the Add/Remove Programs window.
18. Close all open windows. The WINS service is installed. You do not have to restart the machine. You are ready to configure the WINS service.
19. Click the **Start** button on the taskbar, point to **Programs**, point to **Administrative Tools**, click **WINS**, and then click your **server name** in the left pane of the WINS console.

20. Click the **Active Registrations** folder below your server object.
21. Right-click **Active Registrations**. The context menu appears.
22. Click **Find by Owner**.
23. On the **Owners** tab sheet in the Find by Owner dialog box, click **All owners**. Then, click **Find Now**. The list of all WINS records appears in the right pane.
24. Notice all the services listed under the Type column. They are shown with their code (all you saw in the NT WINS Manager) and a description of the service.
25. Close all open windows.

Project 4-5 Testing Name Resolution in Windows 2000



You will need a computer running Windows 2000 that is configured with TCP/IP and on a network with other properly configured TCP/IP computers, including Liverpool. Ask your instructor what is available on the network.

1. If your server is not powered up, power it up now.
2. Press **Control/Alt/Delete**.
3. In the User name box, type **administrator**.
4. In the Password box, type **password**. (If this does not work, ask the instructor for the password.)
5. In the Log on to box, use the selection arrow to select **INTERSALES**. (This, too, will depend on the classroom configuration.)
6. Press **Return**.
7. When the desktop appears, click the **Start** button on the taskbar, and then click **Run**.
8. In the Open box of the Run dialog box, type **cmd**, and then click the **OK** button.
9. At the command prompt, type **ping liverpool** (or a server name given to you by the instructor).
10. If DNS is properly configured (both client and server sides), the name resolution from Liverpool should have been resolved by DNS. If the output resembles the following:

```
Pinging liverpool.intersales.corp [192.168.1.201] with
32 bytes of data
```

then the name resolution was resolved by DNS. If you see simply a server name such as
Liverpool

the name resolution was resolved using a NetBIOS broadcast.

11. Close all open windows on your desktop.



Project 4-6 Configuring a DNS Server to Automatically Update DNS Records

This project requires the completion of Hands-on Project 4-3. In this project, you will configure the DHCP server to automatically update DNS records for DNS clients.

1. If your server is not powered up, power it up now.
2. Press **Control/Alt/Delete**.
3. In the User name box, type **administrator**.
4. In the Password box, type **password**. (If this does not work, ask the instructor for the password.)
5. In the Log on to box, use the selection arrow to select **INTERSALES**. (This, too, will depend on the classroom configuration.)
6. Press **Return**.
7. When the desktop appears, click the **Start** button, point to **Programs**, point to **Administrative Tools**, and then click **DHCP**.
8. Expand the **server name** in the left pane, and then click the **scope** you created in Hands-on Project 4-3.
9. Right-click the **scope**. The context menu for scope appears.
10. Click **Properties**.
11. On the Scope [nnn.nnn.nnn.nnn] dialog box, click the **DNS** tab.
12. On the DNS Properties sheet, make sure the following options are selected:
Automatically update DHCP client information in DNS
Update DNS only if DHCP client requests
Discard forward (name-to-address) lookups when lease expires
Enable update for DNS clients that do not support dynamic update
13. Click the **OK** button.
14. Close all windows on the desktop.

CASE PROJECTS



Case 4-1 A Functional TCP/IP Design

Colorful Paint is a commercial paint manufacturing and distribution company with four sites. Their Chicago headquarters/plant location has 1433 hosts, the Salt Lake City location has 788 hosts, the Houston location has 135 hosts, and the Denver location has 703 hosts.

Colorful has an NT Domain and uses Windows NT on the desktop. The PDC and one BDC are in Chicago; each location has a BDC. Much of the infrastructure has been in place since 1996, except for an upgrade to 10/100 network cards and auto-sensing Ethernet hubs. The existing routers will support up to 220 hosts per subnet.

Users at each location can access network resources at all other locations. The present WAN connections consist of a single connection from each site to the Telco Private Network through which each site can connect to all the other sites. Chicago has a T-1 connection to the WAN and to the Internet. Salt Lake City also has a T-1 connection to the WAN. Denver and Houston each have a fractional T-1 connection of 168 Kbps. All locations access the Internet through a firewall and proxy server at the Chicago location.

For the proposed network, the IT manager wants redundant links for each location.

Private addresses will be used for all internal addresses.

1. Using the information above, recommend a method to provide redundant WAN links. Explain your reasoning. These redundant links are for emergency use and should have a minimum charge to have in place, and usage charges when used.
2. Design an IP addressing implementation, providing the following information: How many public addresses are needed in this design? How many subnets will be within each site? How many subnets with just two hosts may be needed? Using private addresses, design the actual IP usage for this organization.



Case 4-2 Designing for DHCP and DNS

You are a consultant for a company that is contemplating a move from an IPX/SPX Novell 3.12 network (there are some out there) to TCP/IP and Windows 2000 Active Directory. The entire network has been revamped for the first time in eight years, and all hosts have been upgraded to Windows 2000. You have been asked to aid in the DHCP and DNS portion of the design. They have already done their IP addressing planning and will have three subnets with the following number of hosts:

Table 4-3 Planned subnets and hosts

Subnet #	Number of Hosts
1	500
2	100
3	250

All the computers are stationary on the network, located on one campus, and network changes are expected to be very infrequent. They also are rather tolerant of downtime and not willing to spend the money for fault tolerance or on basic infrastructure changes.

Write an outline of your strategy for this design, including placement of DHCP and DNS servers, and your recommendation for which hosts should be DHCP clients.

